



HAL
open science

Graph Subshifts

Pablo Arrighi, Amélia Durbec, Pierre Guillon

► **To cite this version:**

| Pablo Arrighi, Amélia Durbec, Pierre Guillon. Graph Subshifts. 2026. <hal-03238087v4>

HAL Id: hal-03238087

<https://hal.science/hal-03238087v4>

Preprint submitted on 30 Apr 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY-NC-ND 4.0 - Attribution - Non-commercial use - No Derivative Works - International License

Graph subshifts

Pablo Arrighi, Amélie Durbec, and Pierre Guillon

¹ Université Paris-Saclay, Inria, CNRS, LMF, 91190 Gif-sur-Yvette, France
and IXXI, Lyon, France

² Université Paris-Saclay, CNRS, LISN, 91190 Gif-sur-Yvette, France

³ Aix-Marseille Université, CNRS, I2M, Marseille, France

Abstract. We propose a definition of graph subshifts of finite type that can be seen as extending both the notions of subshifts of finite type from classical symbolic dynamics and finitely presented groups from combinatorial group theory. These are sets of graphs that are defined by forbidding finitely many local patterns. In this paper, we focus on the question whether such local conditions can enforce a specific support graph, and thus relate the model to classical symbolic dynamics. We prove that the subshifts that contain only infinite graphs are either aperiodic, or feature no residual finiteness of their period group, yielding non-trivial examples as well as two natural undecidability theorems.

1 Introduction

Subshifts of finite type are well studied objects in symbolic dynamics [17] and ergodic theory. Given an alphabet Σ , a subshift of finite type (SFT) on \mathbb{Z}^d is a set of configurations $\mathcal{Y} \subseteq \Sigma^{\mathbb{Z}^d}$ that do not contain a given finite set of forbidden patterns $\mathcal{F} \subseteq \Sigma^{\{-n, n\}^d}$. In spite of their relatively benign and local definition, SFT have proven to have complex global behaviours. One example of such results is the existence of aperiodic SFTs in dimension 2, but not in dimension 1. Even some natural problems such as determining whether an SFT contains a periodic configuration at all, are proven to be undecidable [5]. This local-to-global complexity is in fact shared by multiple dynamical systems and most notably with Cellular Automata (CA) [16]. However, the two models are deeply connected as CA of dimension d can be seen as subshifts of dimension $d + 1$. This connection has been used to prove multiple theorems on different subclasses of CA [14]. This result is also one of the original motivations of this paper: Graph subshifts ought to be a means to study a generalization of CA called Causal Graph Dynamics (CGD) [1]. CGD extend CA in two complementary ways: first because they are defined on arbitrary graphs of bounded degree m , and second because they allow for the graph itself to evolve, according to a local, shift-invariant rule. Graph subshifts aim to encompass CGD spacetime diagrams as a subclass.

Subshifts have already been generalized to configurations over finitely generated groups. However this is not enough to simulate CGD as all configurations share the same support graph. In this paper, we provide a formalism which relaxes this constraint, allowing the support of the graphs

itself to also be prescribed by set of forbidden patterns. We introduce and formalize the notion of graph subshifts, i.e. sets of graphs that are defined by forbidding local patterns.

The natural question to ask, then, is whether we may enforce, by means of such local constraints only, that the set of graphs be of a particular shape. In particular, notice that a graph subshift may contain a finite or an infinite number of graphs (up to shift), and that the graphs themselves may be of finite or of infinite size. In this first paper we focus on the question whether there exist subshifts whose graphs are all of infinite size. The question is non-trivial; for instance we prove that the problem whether a given set of forbidden patterns induces a graph subshift without finite configuration is undecidable. The problem whether it uniquely fixes the support graph is also shown undecidable. Still, we prove that the graph subshifts that contain only infinite graphs are either aperiodic, or feature no residual finiteness of their period group, yielding us with non-trivial examples and establishing connections with different areas of mathematics: periodicity from symbolic dynamics; residually finite groups from combinatorial group theory [4][18]; and graph covers as used in graph theory and distributed computing.

All proofs are in the Pre-print version of this paper. A very exploratory version of those results has been presented in the (non-proceedings track of) workshop Automata 2021.

2 Graphs

The notions of graphs and of pointed graphs modulo, together with their operations, are intuitive enough as summarized here, for rigorous definitions see [2]. Let π be the finite set of ports, $\Pi = \pi^2$, and V some universe of names.

Graphs. The graphs denoted by letters $G, H \dots$ are the usual, connected, undirected, possibly infinite, bounded-degree labelled graphs, but with a few added twists:

- Vertices are connected through their ports. An edge is an unordered pair $\{x : a, y : b\}$ of $V \times \pi$, or a singleton $\{x : a\}$ standing for a self-loop, where x, y are vertices and $a, b \in \pi$ are ports. Each port is used at most once per node: if $e, e' \in E(G)$ intersect, they must be equal. As a consequence the degree of the graph is bounded by $|\pi|$.
- Each vertex x is given a label $\sigma(x)$, taken within a finite set Σ , also referred to as an internal state. Sometimes the function σ is denoted σ_G to emphasize that it belongs to the description of G .

The set of all (finite and infinite) graphs having ports π , vertex labels Σ is denoted $\mathcal{G}_{\Sigma, \pi}$, or simply \mathcal{G} . The *support* of graph X is simply the structure of vertices linked by edges with ports, forgetting the labeling σ , or equivalently the graph in which the labeling is changed to the monochromatic one.

Pointed graphs modulo. The goal of this paragraph is to define a set of graphs that is a compact space. Indeed, consider a graph G , a vertex $p \in V(G)$, and $D = G_p^r$ the disk of radius r around p , i.e. the subgraph of G induced by the r -neighbors of x in G . Now forgetting about G , this D




Fig. 1: (a) A graph G . (b) A pointed graph $(G, 1)$. (c) A pointed graph modulo X .

is one amongst the many possible disks of radius r that can appear in \mathcal{G} . In order for the set of graphs to form a compact space, one requirement is that the set of disks of radius r be finite. Morally this is the case, since the degree is bounded by $|\pi|$, but vertex names x, y, z jeopardize this; we must get rid of them. Thus,

- Graphs are endowed with a privileged pointed vertex p playing the role of an origin (in particular, there is no empty graph).
- These pointed graphs (G, p) are then considered modulo isomorphism. Notice that isomorphism rename the pointer in the same way as they rename the vertex upon which it points, i.e. $R : (G, p) \mapsto (RG, R(p))$. Thus the relative position of the other vertices w.r.t to the pointer is unchanged.

Pointed graphs modulo are denoted $X, Y \dots$. The set of all pointed (finite and infinite) graphs modulo is denoted $\mathcal{X}_{\Sigma, \pi}$, or simply \mathcal{X} . Given a graph non-modulo G , and a vertex $p \in V(G)$, we write $X((G, p))$ for the corresponding pointed graph modulo. Additionally the set of disks of radius r with states Σ and ports π is written $\mathcal{X}_{\Sigma, \pi}^r$, or simply \mathcal{X}^r .

Paths and vertices. Over graphs modulo isomorphism without pointer, vertices no longer have a unique identifier, which makes designating a vertex a daunting task. Fortunately, our graphs are pointed, so that any vertex \hat{u} can always be designated by some sequence of ports u that leads to it, starting from the origin. Given a pointed graph modulo X , its set of finite paths, starting from the origin, forms a subset $\mathcal{L}(X)$ of Π^* with $\Pi = \pi^2$. Whenever two paths u and v lead to the same vertex, we write $u \equiv_X v$. Thus vertex \hat{u} is really the \equiv_X -equivalence class of u , hence the hat notation. The set of these equivalence classes, a.k.a vertices, is denoted $V(X)$. Notice that, starting from origin, following path $v = (ab) \cdots (cd)$ and then path $\bar{v} = (dc) \cdots (ba)$, leads back to the vertex at the origin, namely $\hat{\varepsilon}$, where ε denotes the empty path.

Remark 1. Notice that two paths $u, u' \in \mathcal{L}(X)$ designate the same vertex $\hat{u} = \hat{u}'$ if and only if for every $v \in \mathcal{L}(X_u)$, $\widehat{uv} = \widehat{u'v}$ (in particular, $v \in \mathcal{L}(X_{u'})$).

Back to graphs non-modulo. Sometimes we still need to manipulate usual graphs, where vertices do have names. For this purpose we use a canonical

naming function $X \in \mathcal{X} \mapsto G(X) \in \mathcal{G}$, which names each vertex of X by the set of paths that lead to it, starting from ε . This $G(X)$ is referred to as ‘the associated graph’. One has $X(G(X), \hat{\varepsilon}) = X$.

Shift. Let $X \in \mathcal{X}$ be a pointed graph modulo. Consider a path $u \in \mathcal{L}(X)$. Then X_u is $X((G(X), \hat{u}))$. The pointed graph modulo X_u is referred to as X *shifted by u* .

From now on we refer to $X \in \mathcal{X}$ as just ‘a graph’, as it is clear from the notation that it is in fact a pointer graph modulo.

3 Subshifts

Given a pointed graph modulo X , a word in Π^* designates a vertex, and a language $L \subseteq \Pi^*$ designates a set of vertices. If, moreover, L is stable under taking the prefix, then the designated vertices induce a connected pointed subgraph of X which we refer to as ‘cut’. Given a set \mathcal{Z} of graphs, one may be interested in selecting those whose cut according to L equals some graph Z , which we refer to as ‘cylinder’.

Definition 1 (Cuts and cylinders). *Consider $L \subseteq \Pi^*$ a prefix-stable language, i.e. such that $uv \in L$ implies $u \in L$. The L -cut of a graph $X \in \mathcal{X}$ is the subgraph induced by the vertices $\{\hat{u} \mid u \in L \cap \mathcal{L}(X)\}$. It is denoted $X|_L$. Consider $\mathcal{Z} \subseteq \mathcal{X}$, we write $\mathcal{Z}|_L$ for $\{X|_L \mid X \in \mathcal{Z}\}$. Consider $Z \in \mathcal{Z}|_L$, the cylinder of Z within \mathcal{Z} is $\{X \in \mathcal{Z} \mid X|_L = Z\}$. It is denoted $[Z]_{\mathcal{Z}}$.*

Definition 2 (Subshift). *Let \mathcal{F} be a set of tuples (F, L) , where each F is a finite graph, and each L a finite prefix-stable language. The subshift forbidding \mathcal{F} is*

$$\mathcal{Z} = \{X \in \mathcal{X}_{\Sigma, \pi} \mid \forall v \in \mathcal{L}(X), \forall (F, L) \in \mathcal{F}, (X_v)|_L \neq F\}.$$

It is of finite type if \mathcal{F} can be chosen finite.

Remark 2. In the pairs (F, L) , L cannot always be assumed to be the language labelling the paths of F . For example, requiring every vertex to have an edge with port $a \in \pi$ is a very natural finite-type condition, but this SFT is not of the form $\{X \in \mathcal{X}_{\Sigma, \pi} \mid \forall v \in \mathcal{L}(X), \forall F \in \mathcal{F}, (X_v)|_{L_F} \neq F\}$ as the L_F do not contain a and hence cannot tell that it is absent.

Yet, this SFT can be defined through our definition by

$$\{(F, \{\varepsilon, ap\}) \mid p \in \pi, F \text{ 1-node graph}\}.$$

The same remark holds for the dual definitions, *i.e.* through allowed patterns: think of the SFT consisting of those graphs having no port $a \in \pi$.

Lemma 1. *The following are equivalent.*

1. \mathcal{Z} is a subshift of finite type.
2. \mathcal{Z} is the subshift forbidding some $\mathcal{F}' \times \{M\}$ where \mathcal{F}' is a finite set of finite graphs and M a single finite prefix-stable language.

3. \mathcal{Z} is the set of graphs allowing some finite set \mathcal{A} of tuples (A, L) , where A is a finite graph and L a finite prefix-stable language—in the sense that:

$$\mathcal{Z} = \{ X \in \mathcal{X}_{\Sigma, \pi} \mid \forall v \in \mathcal{L}(X), \exists (A, L) \in \mathcal{A}, (X_v)_{|L} = A \}.$$

4. \mathcal{Z} is the set of graphs allowing $\mathcal{A}' \times \{M\}$ where \mathcal{A}' is a finite set of finite graphs and M a single finite prefix-stable language.

Definition 3 (Defining window). *The prefix-stable language M is the same in Conditions 2 and 4 of Lemma 1. It is called a defining window for \mathcal{Z} .*

Remark 3. The duality between forbidding and allowing no longer holds in general when the subshifts are not of finite type. For example we can think about the set of graphs authorizing exactly those finite graphs which have exactly one red vertex. This set however is not closed topologically (because the red vertex can in some sense vanish to infinity; cf. [2] as summarized by the Appendix) and hence it is not a subshift. Its closure is the so-called *one-dimensional sunny-side-up* subshift, as defined by forbidding those finite graphs having two red vertices.

Remark 4. Given any subshift \mathcal{Z} , one can define the subshift of its infinite graphs, by additionally forbidding $(F, L\pi^2)$ for every prefix-stable language $L \in \Pi^*$ and every graph F such that $F_{|L} = F$. Indeed this condition exactly means that L was already describing all possible vertices in F , and longer paths do not add any. Nevertheless, this subshift is very rarely of finite type, as we will see in the following section.

On the contrary, all SFTs presented so far contain finite graphs. It is not obvious a priori whether some SFTs contain only finite graphs, or even whether they are homogeneous, in the sense that all the graphs share a common support. We will discuss these properties in the next sections.

4 Cayley SFTs

The equivalent definitions for SFTs all capture the concept of local constraint, and in the following sections, we often define our SFTs by describing some conditions that graphs must locally satisfy. Among these conditions, *nearest-neighbor constraints* are those defined with a single edge between two vertices, that is $L \subseteq \{\varepsilon\} \sqcup \Pi$.

Example 1. A *directed 2-regular SFT* is an SFT with ports $\pi = \{a, a', b, b'\}$ obtained by imposing nearest-neighbor constraints:

1. enforcing that edges be of the form pp' ;
2. enforcing that all ports be occupied.

It contains, amongst many others, the Cayley graph of the free group over two generators (a tree), as well as those of all of its quotients.

Example 2 (Locally grid-like SFT). We can define a graph SFT with ports $\pi = \{a, a', b, b'\}$ corresponding to the monochromatic fullshift on $\langle a, b \mid aba^{-1}b^{-1} = \varepsilon \rangle$, by imposing three families of constraints:

1. enforcing that edges be of the form pp' ;
2. enforcing that all ports be occupied;
3. enforcing that the subgraph induced by $\{\varepsilon, aa', bb', aa'bb'\}$ obeys the relation *i.e.* form a square as in Fig. 2).

This graph SFT obviously contains the Cayley graph of $\langle a, b | aba^{-1}b^{-1} = \varepsilon \rangle$, but also the Cayley graphs of all quotient groups (excluding really small groups such as the trivial group).




Fig. 2: *Locally grid-like SFT corresponding to the Cayley graph of $\langle a, b | aba^{-1}b^{-1} \rangle$. The SFT is defined by forbidding all but these patterns, on their corresponding languages.*

A similar construction can be performed for any finitely presented group, and via a more precise third constraint, get a cleaner statement (with all possible quotients).

Definition 4 (Cayley SFT). *Consider a finitely presented group $\Gamma = \langle \mathcal{J} | \mathcal{R} \rangle$, that is, a group generated by the abstract finite set \mathcal{J} , where \mathcal{R} represents a set of relations, that is words over $\mathcal{J} \sqcup \mathcal{J}^{-1}$ that represent the identity element (and from which all such words can be derived). For example $\mathbb{Z}^2 = \langle a, b | aba^{-1}b^{-1} \rangle$ (it will always be implicitly given by this canonical presentation) and $\mathbb{F}_2 = \langle a, b \rangle$.*

For every alphabet Σ , there is a canonical way to associate a Cayley SFT $\mathcal{Z}_\Sigma^\Gamma$ to the group Γ (in fact, to the pair (Γ, \mathcal{J}) , which is a marked group when \mathcal{J} is symmetric and does not contain the identity ε). Define $\pi = \mathcal{J} \sqcup \mathcal{J}^{-1}$ and the following constraints:

1. *each edge involves two inverse ports;*
2. *every vertex must have all possible ports;*
3. *for every word $u \in \mathcal{R}$ and L its language of prefixes, one forbids the pair (Y, L) for every graph Y supported by L in which path u is not a cycle back to the origin.*

Let Γ be a finitely presented group. A SFT over alphabet Σ is defined by Γ -NN constraints if they include precisely the forbidden patterns defining

$\mathcal{Z}_\Sigma^\Gamma$ and some additional nearest-neighbor constraints. Note that if $\Gamma = \mathbb{F}_k$, then \mathbb{F}_k -NN constraints are simply constraints on the edges, provided the correct regularity of the graph.

Remark 5. The Cayley SFT obviously contains the Cayley graph of Γ , but also all Cayley graphs of its quotient groups, as will be seen in the next section. For example, $\mathcal{Z}_\Sigma^{\mathbb{Z}^2}$ contains the infinite grid as well as tori and cylinders, and $\mathcal{Z}_\Sigma^{\mathbb{F}_2^2}$ contains Cayley graphs of all groups with two generators.

Any classical nearest-neighbor SFT defined over some Cayley graph Γ is included in the corresponding graph SFT defined by Γ -NN constraints. To define a subshift whose unique support is the Cayley graph of Γ , one generally needs to add many constraints imposing that paths that are not labeled by words representing the identity element be not cycles. This usually yields a subshift of infinite type.

Example 3 (Hard-square model). The hard-square model SFT is defined by \mathbb{Z}^2 -NN constraints: $\mathcal{Z} \subseteq \mathcal{X}_{\{0,1\},\{a,a',b,b'\}}$ such that for every configuration X and edge $\{x : a, y : b\} \in E(X)$, one has that $\sigma_X(x) = 0$ or $\sigma_X(y) = 0$. This can be expressed as an SFT by forbidding the 0/1-coloured versions of Fig. 2 that have two adjacent 1-coloured vertices. This is a direct generalization of a well-known SFT over \mathbb{Z}^2 . As a graph SFT, it not only contains all the configurations that are already present on \mathbb{Z}^2 such as in Fig. 3, but also new infinite configurations such as the cylinder of Fig. 4, and even finite configuration such as tori only containing 0.




Fig. 3: Configuration of the hard-square model that is present in \mathbb{Z}^2 . The colours represent the 0/1 internal states of the vertices.

Besides finitely presented groups, the notion of quotient of a group can be extended to periodic graphs, which is the subject of the following section.




Fig. 4: Configuration of the hard-square model whose support graph is not \mathbb{Z}^2 . In order to preserve clarity, only half of the ports are written and some edges have been dashed.

5 Periodicity and finiteness

5.1 Quotient and Periods

Definition 5 (Period). Any word $u \in \mathcal{L}(X)$ such that $X = X_u$ is called a path-period. The corresponding vertex \hat{u} is called a period. The stabilizer of X is its set of periods; it is denoted by $\mathcal{S}(X)$. X is called weakly periodic if $\mathcal{S}(X) \supsetneq \{\hat{\varepsilon}\}$, strongly nonperiodic if $\mathcal{S}(X) = \{\hat{\varepsilon}\}$.

Notice that if X is weakly periodic, then $\mathcal{S}(X)$ is infinite or every period \hat{u} is torsion, i.e. $\widehat{u^n} = \hat{\varepsilon}$ for some $n \geq 2$.

Notice that monochromatic Cayley graphs are exactly the graphs whose all vertices are periods. As we have seen in our Cayley SFT examples, periodic graphs can typically be “folded into smaller graphs” that also belong to the graph SFT. The following formalizes the idea of a graph folding into another one.

Definition 6 (Homomorphism). A homomorphism from graph $X \in \mathcal{X}$ to graph $X' \in \mathcal{X}$ is a function φ from $V(X)$ to $V(X')$ such that for all $u \in \mathcal{L}(X)$,

$$\varphi(\hat{u}) = \hat{u} \quad \text{and} \quad \sigma_X(\hat{u}) = \sigma_{X'}(\hat{u}).$$

Note the first equality implies $\mathcal{L}(X) \subseteq \mathcal{L}(X')$. When $\mathcal{L}(X) = \mathcal{L}(X')$ (note that this implies that φ is surjective and locally injective), we then say that X covers X' and that φ is a covering.

Remark 6. A homomorphism sends cycles to cycles and does not change the set of path-periods (i.e. if u' is such that $\hat{u}' = \phi(\hat{u})$, then $X = X_u \iff X' = X_{u'}$). Hence if \hat{u} is a period of X , $\varphi(\hat{u})$ is a period of $\varphi(X)$.

Definition 7 (Quotient graph). Consider $X \in \mathcal{X}$ and a stabilizer subgroup H . We can define a quotient graph X/H to be $X((G, p))$ where:

- $V(G) = \{H\hat{u} \mid \hat{u} \in V(X)\}$, i.e. each vertex is labelled by a set $H\hat{u} = \{Hv \mid v \equiv_X u\}$.
- $p = H$, i.e. the pointer is the vertex labeled by H .
- $E(G) = \left\{ (H\hat{u} : a, H\widehat{uab} : b) \mid uab \in \mathcal{L}(X) \right\}$.

$$- \sigma_G(H\hat{u}) = \sigma_X(\hat{u}).$$

By stabilizer subgroup, we mean a subgroup of the stabilizer group. As the elements of a stabilizer subgroup cannot be distinguished modulo isomorphism, it follows that a stabilizer subgroup induces a natural covering:

Proposition 1. *For every graph $X \in \mathcal{X}$ and every stabilizer subgroup H of X , X/H is a well-defined graph, and $\phi : \hat{u} \mapsto H\hat{u}$ is a covering from $V(X)$ to $V(X/H)$.*

On top of being able to "fold" a graph, we can do so while locally preserving it, given that a certain condition on the stabilizer subgroup is met:

Definition 8 (Separated). *If $U, V \subset \Pi^*$, then we denote $UV := \{uv \mid u \in U, v \in V\}$, and $\bar{U} := \{\bar{u} \mid u \in U\}$. If $M \subseteq \Pi^*$ is a prefix-stable language, we say that a set H of vertices is M -separated in graph X if for every $\hat{u} \in V(X)$, there at most one $\hat{v} \in H$ such that there exists $w \in M\bar{M}$ with $\widehat{wv} = \hat{v}$.*

Proposition 2. *Let X be a graph, $M \subseteq \Pi^*$ a prefix-stable language, and $H \leq \mathcal{S}(X)$ a stabilizer subgroup which is M -separated in X . Then the covering preserves the language of support M in the sense that for all $u \in \mathcal{L}(X)$, $(X_u)_{|M} = (X/H_u)_{|M}$ (modulo isomorphism). In particular, if X belongs to some SFT \mathcal{X} with defining window M , then X/H also belongs to \mathcal{X} .*

5.2 Strong periodicity

Definition 9 (Density, strong periodicity). *Consider a graph X and a finite prefix-stable language $M \subset \Pi^*$. We say that a language $L \in \Pi^*$ is M -dense in X if for all $\hat{u} \in V(X)$, $uM \cap L \neq \emptyset$. L is said dense in X if such a M exists. A graph X is said to be strongly periodic if its stabilizer $\mathcal{S}(X)$ is dense in X . Equivalently, there are only finitely many distinct (modulo isomorphisms) graphs X_u , for $u \in \mathcal{L}(X)$.*

Note that finite graphs are strongly periodic, but may not always be weakly periodic. Apart from this degenerate case, strongly periodic graphs are weakly periodic.

Notice also that strong periodicity is almost equivalent to being a Cayley graph, up to replacement of each vertex by a fixed finite graph (with possible edges between them). This replacement (factor map) corresponds somehow to the application of a causal graph dynamics [2] to the Cayley graph.

Lemma 2. *A graph X is strongly periodic if and only if X/H is finite for some stabilizer subgroup H .*

The following result establishes a connection between the SFT that admit finite graphs, and the well-established notion (see for instance [7, Chap. 2]) of residual finiteness from group theory:

Definition 10 (Residually finite group). *A group \mathcal{S} is residually finite if and only if for any non-identity element $\hat{u} \in \mathcal{S}$, there exists a normal subgroup $H \trianglelefteq \mathcal{S}$ of finite index such that $\hat{u} \notin H$.*

This notion is very robust. In particular, the property extends to finite sets of non-identity elements I , *i.e.* there exists a normal subgroup $H \trianglelefteq \mathcal{S}$ of finite index such that $H \cap I = \emptyset$, simply because any finite intersection of finite-index subgroups is a finite-index subgroup.

Let us relate residual finiteness to quotient graphs, at the intuitive level first. Say X belongs to some SFT \mathcal{Y} and let $\hat{u}_1, \dots, \hat{u}_n$ be distinct elements of $V(X)$. The above property gives a subgroup H of its stabilizer, and thus X/H is a possible quotient graph; one which maintains that $\hat{u}_1, \dots, \hat{u}_n$ are distinct. The correct choice of $\hat{u}_1, \dots, \hat{u}_n$ will help us make sure that X/H still belongs to SFT \mathcal{Y} .

Theorem 1. *Let \mathcal{Y} be a SFT. \mathcal{Y} contains a finite graph if and only if \mathcal{Y} contains a strongly periodic graph X such that the stabilizer $\mathcal{S}(X)$ is residually finite.*

Thus there are two complementary ways in which a nonempty SFT may have only infinite configurations: either by preventing strongly periodic configurations, or by forcing them to have stabilizers that are not residually finite.

In the latter case (for example when the stabilizer is a Tarski monster group, see for instance [21]), one can build SFTs with strongly periodic configurations that do not fold uniformly onto finite graphs, as shown by the following.

Remark 7. Γ is not residually finite if and only if $\mathcal{Z}_{\{0\}}^\Gamma$ includes an SFT containing the Cayley graph of Γ but no finite quotient of it. Indeed, one can consider the SFT \mathcal{Z} included in $\mathcal{Z}_{\{0\}}^\Gamma$ defined by the additional constraint that u is never a cycle, where u is a noncyclic path that witnesses that Γ is not residually finite, in the sense that all finite quotients of Γ map u to ε .

5.3 Quotients in Cayley SFTs

We have seen that the notions of quotients and SFTs are linked, because finite-type constraints cannot easily tell a graph from some of its quotients. On the opposite point of view, it is noticeable that some SFTs admit a maximal support with respect to coverings.

Proposition 3. *Let Γ be a finitely presented group, X a graph, and \mathcal{Z} some SFT defined by Γ -NN constraints. The following are equivalent:*

1. $X \in \mathcal{Z}$.
2. X is covered by some colored Cayley graph Y of Γ which is in \mathcal{Z} .

Remark 8. In graphs of $\mathcal{Z}_\Sigma^\Gamma$, all stabilizers are subgroups of Γ (because the concatenation law is coherent with that of the group). One has to be careful though, that quotients of Cayley graphs may not be Cayley graphs, because the stabilizer is not always a normal subgroup. Nevertheless, it is

classical that every finite-index subgroup includes a finite-index normal subgroup (see for instance [7, Lemma 2.1.10]), so that a graph of $\mathcal{Z}_{\Sigma}^{\Gamma}$ is strongly periodic if and only if some stabilizer subgroup is a finite-index normal subgroup of Γ .

In particular, in graphs of $\mathcal{Z}_{\Sigma}^{\mathbb{Z}^2}$, all stabilizers are subgroups of \mathbb{Z}^2 , which makes many properties easier.

Proposition 4. *Let \mathcal{Z} be an SFT defined by \mathbb{Z}^2 -NN constraints. Then exactly one of the following occurs:*

1. *Every graph in \mathcal{Z} is supported by the infinite grid, and is strongly nonperiodic.*
2. *\mathcal{Z} contains a finite graph and a strongly periodic infinite grid.*

It is already interesting to note that the dichotomy from Proposition 4 is nontrivial, directly from the existence of strongly aperiodic SFTs over the grid (and other ones). But the next subsection will show a stronger result.

5.4 Undecidability

Theorem 2. *The dichotomy from Proposition 4 is undecidable, given a list of \mathbb{Z}^2 -NN constraints.*

This is a consequence of Theorem 1, allowing us to reduce the previous dichotomy to a classical periodicity problem on \mathbb{Z}^2 . Consequently, the class of graph SFTs that contains only infinite graphs with the same support is computably unseparable from the class of graph SFTs containing both finite and infinite graphs.

This means the undecidability of any property that is implied by one of the two classes, and which implies the negation of the other one.

Corollary 1. *The problem whether a graph SFT contains only infinite graphs, and the problem whether a graph SFT admits only one support for all its graphs, are both undecidable (even when restricted to graph SFTs with \mathbb{Z}^2 -NN constraints).*

More precisely, the reduction directly gives that all properties separating the dichotomy are Σ_1^0 -hard. The encoding of Cayley graphs into graph SFTs also, more directly, implies the following undecidability result.

Theorem 3. *There exists a finite presentation of a group Γ such that the problem whether a given word $u \in \Pi^*$ represents a cycle from the origin back to itself in every graph of $\mathcal{Z}_{\{0\}}^{\Gamma}$ is undecidable.*

It is actually Σ_1^0 -complete.

6 Conclusion

To sum up, we provided a formalism generalizing SFTs on Cayley graphs by allowing the support graph to be determined by forbidden subgraphs. The generalization also allows the SFT to have finite configurations and we provided examples of such SFT. By defining quotient graphs, we have shown that the only way for an SFT to not have finite configurations is to only have configurations that are either nonperiodic, or with a stabilizer that is not residually finite. Similarly, we have shown that having only one support graph requires configurations to have finite stabilizers. We used both results to prove that the finite configuration problem and the support graph unicity problem are undecidable.

In the future, we wish to study determinism and the relation between causal graph dynamics and graph subshifts, in the perspective of proving structure and complexity results on causal graph dynamics. Moreover, graph subshifts could provide a new point of view for studying group subshifts. On the other hand, many notions and results known over group subshifts such as sofic subshifts, determinism or entropy, would be interesting to lift to graph subshifts.

We have seen that graph subshifts could generalize both classical symbolic dynamics and finite-type group geometry. They could be used to simulate a larger class of objects; let us discuss one example. Let \mathcal{T} be a finite set of *tiles*, *i.e.* polygons of the plane \mathbb{R}^2 . A *tiling* is then a covering of \mathbb{R}^2 by translated copies of tiles from \mathcal{T} , the intersection of any two of them being either empty, or a vertex, or a whole edge, or the whole tile. Now let us define subshift $\mathcal{Z}_{\mathcal{T}}$ with $\Sigma = \mathcal{T}$, π as the set of tile edges up to translation, and the following constraints: each vertex labeled $t \in \mathcal{T}$ must have exactly those ports which correspond to the edges of t , and every simple cycle must correspond to a valid vertex in a tiling of the plane by \mathcal{T} (*i.e.* it involves tiles that can be put next to each other around this vertex). The fact that the latter condition can be forced by finitely many constraints (so that $\mathcal{Z}_{\mathcal{T}}$ be an SFT) is equivalent to the classical *finite-local-complexity* property of \mathcal{T} . With this construction, there is a one-to-one correspondence between tilings by \mathcal{T} of the plane and of all surfaces and graphs in $\mathcal{Z}_{\mathcal{T}}$. Now if one is interested in the shift-orbit closure of a single tiling of the plane (*i.e.* the Penrose tiling), it is an interesting question whether this corresponds to an SFT in $\mathcal{Z}_{\mathcal{T}}$. Some work has been devoted to understanding for which tilings one can force locally tilings of the plane to be in the shift-orbit closure [9,6]; one says that the tilings admits *local rules*. Additionnally forbidding (via finite constraints) all possible graphs which represent tilings on surfaces but not on the plane is a new question. Our work shows that if the tiles are rectangles and if we allow some decorations (that could also be encoded locally as small hanging subgraphs), then this is possible. It is tempting to think that this remains possible in the case of general strongly nonperiodic tilings, but the proofs would require a new formalism, because the support is no longer a Cayley graph. The same could be done in the tridimensional space, or in the hyperbolic plane.

Let us make a final remark. Graph theorists are used to considering families of graphs defined by forbidden finite graph minors. Of course,

seen as including possibly infinite graphs, these families are subshifts, but they are a priori not SFTs (not even sofic). However, even if the *minor* operation involves vertices at possibly unbounded distance, it is known to have some kind of representation as local constraints [15] (involving a nondeterministic coloring). The relation with our setting could be investigated.

Acknowledgements

We thank the referees for their careful reading of the preliminary version, and helpful remarks. We are indebted to Nicolas Schabanel who encouraged us to study graph subshifts and contributed to the first four definitions. This publication was made possible through the support of the ID# 62312 grant from the John Templeton Foundation, as part of the ‘The Quantum Information Structure of Spacetime’ Project (QISS). The opinions expressed in this publication are those of the author(s) and do not necessarily reflect the views of the John Templeton Foundation.

References

1. Pablo Arrighi, Amélia Durbec, and Aurélien Emmanuel. Reversibility vs local creation/destruction. In Michael Kirkedal Thomsen and Mathias Soeken, editors, *Reversible Computation - 11th International Conference, RC 2019, Lausanne, Switzerland, June 24-25, 2019, Proceedings*, volume 11497 of *Lecture Notes in Computer Science*, pages 51–66. Springer, 2019. doi:10.1007/978-3-030-21500-2_4.
2. Pablo Arrighi, Simon Martiel, and Vincent Nesme. Cellular automata over generalized cayley graphs. *Mathematical Structures in Computer Science*, 28(3):340–383, 2018. doi:10.1017/S0960129517000044.
3. Alexis Ballier. *Propriétés structurelles, combinatoires et logiques des pavages*. Ph.d. thesis, Université de Provence, Nov 2009.
4. Gilbert Baumslag and Donald Solitar. Some two-generator one-relator non-hopfian groups. *Bulletin of the American Mathematical Society*, 68(3):199–201, 1962. 00471. URL: <https://www.ams.org/bull/1962-68-03/S0002-9904-1962-10745-9/>, doi:10.1090/S0002-9904-1962-10745-9.
5. Robert Berger. *The undecidability of the domino problem*. Phd, Harvard University, jul 1964.
6. Nicolas Bédaride and Thomas Fernique. When periodicities enforce aperiodicity. *Communications in Mathematical Physics*, 335(3):1099–1120, May 2015. arXiv: 1309.3686. URL: <http://arxiv.org/abs/1309.3686>, doi:10.1007/s00220-015-2334-8.
7. Tullio Ceccherini-Silberstein and Michel Coornaert. *Cellular Automata and Groups*. Springer, 2009.
8. Benjamin Hellouin de Menibus. Addendum to: Tiling an arbitrary amenable group is decidable. draft, 2022.
9. Thomas Fernique and Mathieu Sablik. Weak colored local rules for planar tilings. *arXiv:1603.09485 [cs, math]*, March 2016. URL: <http://arxiv.org/abs/1603.09485>.

10. Jean-Marc Gambaudo. A note on tilings and translation surfaces. *Ergodic Theory & Dynamical Systems*, 26(1):179–188, 2006.
11. Yuri Gurevich and I. Koryakov. A remark on berger’s paper on the domino problem. *Siberian Journal of Mathematics*, 13:459–463, 1972.
12. G. A. Hedlund. Endomorphisms and automorphisms of the shift dynamical system. *Math. Systems Theory*, 3:320–375, 1969.
13. Emmanuel Jeandel and Pascal Vanier. The Undecidability of the Domino Problem. In Shigeki Akiyama and Pierre Arnoux, editors, *Substitution and Tiling Dynamics: Introduction to Self-inducing Structures*, Lecture Notes in Mathematics, page 66. Springer, 2017.
14. Jarkko Kari. The nilpotency problem of one-dimensional cellular automata. *siam Journal on Computing*, 21(3):571–586, 1992.
15. Dietrich Kuske and Markus Lohrey. Logical aspects of Cayley-graphs: the group case. *Annals of Pure and Applied Logic*, 131(1-3):263–286, 2005. URL: <http://linkinghub.elsevier.com/retrieve/pii/S0168007204000831>, doi:10.1016/j.apal.2004.06.002.
16. C. G. Langton. Self-reproduction in cellular automata. *Physica D*, 10:135–144, 1984.
17. Douglas A. Lind and Brian Marcus. *An introduction to symbolic dynamics and coding*. Cambridge University Press, Cambridge ; New York, 1995. 03085.
18. Roger C. Lyndon and Paul E. Schupp. *Combinatorial Group Theory*. Classics in Mathematics. Springer-Verlag, Berlin Heidelberg, 2001. 04184. URL: <https://www.springer.com/gp/book/9783540411581>.
19. Thierry Monteil. Down with hierarchy !, 2013. FrAC.
20. P. S. Novikov. On the algorithmic unsolvability of the word problem in group theory. *Proceedings of the Steklov Institute of Mathematics*, 44:1–143, 1955. in Russian.
21. A Yu Ol’shanskiĭ. *Geometry of defining relations in groups*, volume 70. Springer Science & Business Media, 2012.
22. Steven Piantadosi. Symbolic dynamics on free groups. *Discrete & Continuous Dynamical Systems - A*, 20(3):725–738, 2008. URL: <http://aimsciences.org//article/doi/10.3934/dcds.2008.20.725>, doi:10.3934/dcds.2008.20.725.